

Before the
Federal Communication Commission
Washington, D.C. 20554

In the Matter of)	
)	
Notice of Proposed Rulemaking,)	FCC 06-103
Notice of Inquiry,)	
and Order)	

ET Docket Nos. 06-135, 05-213, 03-92, and RM-11271

REPLY COMMENTS OF INTEL CORPORATION

October 31, 2006

SUMMARY

The Commission has proposed establishing a new service for advanced medical radio communication (“MedRadio”) devices in the 401-406 MHz band. Intel Corporation (“Intel”) believes that a MedRadio band for medical devices would foster a new ecosystem of personal medical devices that could greatly improve the quality of life for many patients. These devices will provide home and mobile monitoring of chronic diseases, cognitive decline disorders, post operative care, infant care, as well as many other general health and wellness monitoring applications. Our analysis of use case scenarios and technical factors strongly supports the Commission’s proposal. Intel has also provided responses to several of the Commission’s requests for additional information regarding the allocation of spectrum for medical devices.

Before the
Federal Communication Commission
Washington, D.C. 20554

In the Matter of)	
)	
Notice of Proposed Rulemaking,)	FCC 06-103
Notice of Inquiry,)	
and Order)	

ET Docket Nos. 06-135, 05-213, 03-92, and RM-11271

REPLY COMMENTS OF INTEL CORPORATION

Intel Corporation (“Intel”) hereby submits the following reply comments in response to the public notice released by the Federal Communications Commission (“Commission”) in the above referenced proceeding. Intel is the world’s largest semiconductor manufacturer and is a leader in standards and technical innovation. Of particular relevance here, Intel’s Digital Health Group focuses specifically on technologies, devices, standards and services for the healthcare industry.

This corporate division of Intel has reviewed a wide range of consumer healthcare devices, services, and usage scenarios--many of which include the use of wireless sensors. These wireless sensors can be deployed throughout a home or attached to an individual to monitor vital signs, detect severe health traumas or events, and potentially activate life critical responses to conditions. The advantages of wireless technologies in these devices include: greater patient comfort which increases compliance with a monitoring or treatment program and the ease of installation for stationary sensors distributed through a living environment.

However, the use of wireless technologies must address two well recognized challenges; (1) overall power consumption must be reduced to maximize battery life and minimize battery size and (2) the quality of service or overall reliability of the data transmission must be kept high. While some health/wellness companies can select low power wireless technologies that operate in the unlicensed spectrum, depending on the environment these bands can compromise the quality of the data transmission. Additionally, many medical sensors cannot bear the risks associated with operating in unlicensed spectrum.

Accordingly, Intel supports the Commission's proposal to establish a "MedRadio" service. Additionally, Intel believes that the Commission should

define the types of devices that may use the MedRadio service in order to minimize complications that could arise if this spectrum was exploited for non-medical usage. Intel proposes that the MedRadio the band of spectrum should only be used by sensors that collect and transmit physiological data that is used for monitoring or treatment of a patient. These devices may be operated by the patient or their professional healthcare provider. It should be acknowledged that the sensors that use MedRadio and the data that is transmitted requires a high quality of service due to the severity of the medical condition being monitored or treated. It would be inappropriate to allow the MedRadio spectrum to be used by sensor devices that may arguably transmit physiological data, but are used for general fitness where the necessity of the data transmission is not life critical. For all other sensor devices that may be used in health and wellness usage scenarios, we believe that unlicensed spectrum should be used and that there are several robust radio standards that will adequately address these sensor requirements. This would include fitness devices such as pedometers and treadmills as well as home security or fall detection sensors that might be used for an elderly monitoring service.

The Commission has also asked for comments in regards to the usage of standards for the various medical radio services being discussed. Intel has a history in the development and usage of communication standards for a

wide range of industries and we believe that communication standards are a necessary component to allow free and open collaboration within a marketplace between the various companies that develop devices and services. For that purpose, Intel and several other companies recently formed an industry organization called the Continua Health Alliance¹ to specifically enable an ecosystem of interoperable devices for consumers and organizations to better manage their health and wellness. This Alliance is comprised of a wide variety of companies including: medical device vendors, medical implant vendors, fitness device vendors, healthcare providers, cellular vendors, fitness service providers, consumer electronic companies, and technology component companies. All of these companies have joined this Alliance as they believe that through the efforts of a collaborative industry organization, they can enable a personal telehealth eco-system where many diverse vendors can combine their products into new value propositions with significant health benefits for people worldwide. The Continua Health Alliance will be selecting international communication and data standards for different medical and fitness device categories and will be developing guidelines on how vendors and service providers can use these standards to achieve strict interoperability. The Alliance will also create and administer an interoperability testing and certification program whereby vendors will be able use a recognizable certification mark on their device after passing an interoperability test process.

¹ <http://www.continuaalliance.org>

Finally, Intel largely agrees with the Commission's various technical proposals. Appendices 1-4 contain supporting analysis for the below conclusions:

1. We found that even with 25uW of power, significant range can be established. However this could be justified as body worn medical devices may have low efficiency antenna designs. Such low efficiency antennas may require a higher RF field strength which could be traded off with available range and bandwidth. But, without data about the body unit antenna efficiency and noise figure, it is difficult to ascertain the amount of field strength needed.
2. Proportionally decreasing the bandwidth and data rate, for example to 100 KHz with 100 kbps, could potentially increase the available Eb/No (SNR per bit) provided that the device is transmitting a reasonable number of bits/sec/Hertz. However, decreasing the bandwidth and data rate too much may not allow for interesting bandwidth intensive use case scenarios to occur. Please see Appendix 1 (SNR Trade off analysis).

3. We believe that Commission could support allowing both LBT and non-LBT² modes of operation. Our analysis showed that LBT was slightly better for this radio band but with added complexity. It was also found that relying on physical separation alone to support multiple users is not practical and even a simple protocol would greatly increase the user density. It would be best for devices that use LBT to have ACK/NACK mechanisms for reliable transfer especially to compensate for lower SNR or longer distances. One important aspect is that LBT devices and non-LBT devices be in separate bands. In theory, the non-LBT device will always talk without listening and the LBT device will not talk unless the frequency is quiet. Hence, if LBT devices are mixed with non-LBT devices, the result could be that the LBT devices seldom get a chance to make a transmission. Placing the non-LBT devices in the wing bands seems a reasonable thing to do. Intel has done some preliminary analysis to show why it agrees with Commission. Please see Appendix 2 (Multiuser Density analysis - LBT Vs Non-LBT analysis).
4. For bandwidth intensive devices, decreased duty cycle could be achieved thus improving the overall deployment density. Please see Appendix 3 (Trading-off Increased BW for Reduced Duty Cycle)

² It appears to Intel that non-LBT operation is similar to the Aloha protocol that was developed in the mid-1970's for early packet radio. ref. http://en.wikipedia.org/wiki/Aloha_protocol#The_ALOHA_protocol

5. While we support the reduced transmission power level of 250 nW for non-LBT devices, we also believe satisfactory performance could be achieved for higher transmission power levels up to the allowed limit of 25 uW with a duty cycle of 0.1%. Please see Appendix 4 (Decreased transmit power, duty cycle analysis).

Respectfully submitted,

By: /s David L Whitlinger

/s Nandakishore

Kushalnagar

David L. Whitlinger
Kushalnagar
Director, Healthcare Standards Development
Digital Health Group
Group
Intel Corporation

Nandakishore
Wireless Architect
Digital Health

Intel Corporation

/s Richard D Roberts

Richard D. Roberts
Regulatory Research Scientist
Corporate Technology Group
Intel Corporation

Appendix 1 (SNR Trade Off Analysis)

Surprisingly, even at 25 uW of TX power, the theoretical free space range of a MedRadio is large. The plot below shows free space range vs. data rate.



Figure 1 – Free Space Range vs. Data Rate

Assumptions:

- TX power: 25 uW
- Data Rate: variable
- TX ant: 0 dBi
- RX ant: 0 dBi
- Frequency: 403.5 MHz
- Required Eb/No: 10 dB

- Implementation Loss: 3 dB
- RX Noise Figure: 6 dB
- Excess Propagation Loss: 0 dB

This amount of range may be somewhat surprising but this is due to the careful selection of the operating frequency. To give a feel for why the large range, consider that 25 uW of power is -16 dBm. At 1 meter, the aperture loss is ~24 dB so the TX signal at 1 meter is ~ -40 dBm. On the other hand, for an assumed RX 3dB bandwidth of 225 kHz and an assumed modulation efficiency of 1 bps/Hz, the noise power is at ~ -110 dBm. Thus at 1 meter the TX signal is 70 dB above the noise floor. Assuming we need a 10 dB SNR for a low bit error rate, this leaves ~60 dB of SNR that can be traded for range. The relationship is $20 \cdot \log_{10}(\text{distance})$ which means 60 dB can be traded for a 1000 times increase in distance which takes us out to ~1km. For a data rate of 225 kbps, the actual distance for the stated assumptions is 1.1 km.

The next plot shows E_b/N_0 (signal to noise ratio per bit) vs. distance for a fixed data rate of 225 kbps. As the range decreases, the available SNR correspondingly increases. This increased SNR with decreased range can be traded-off for high user density. This is discussed later in this response.

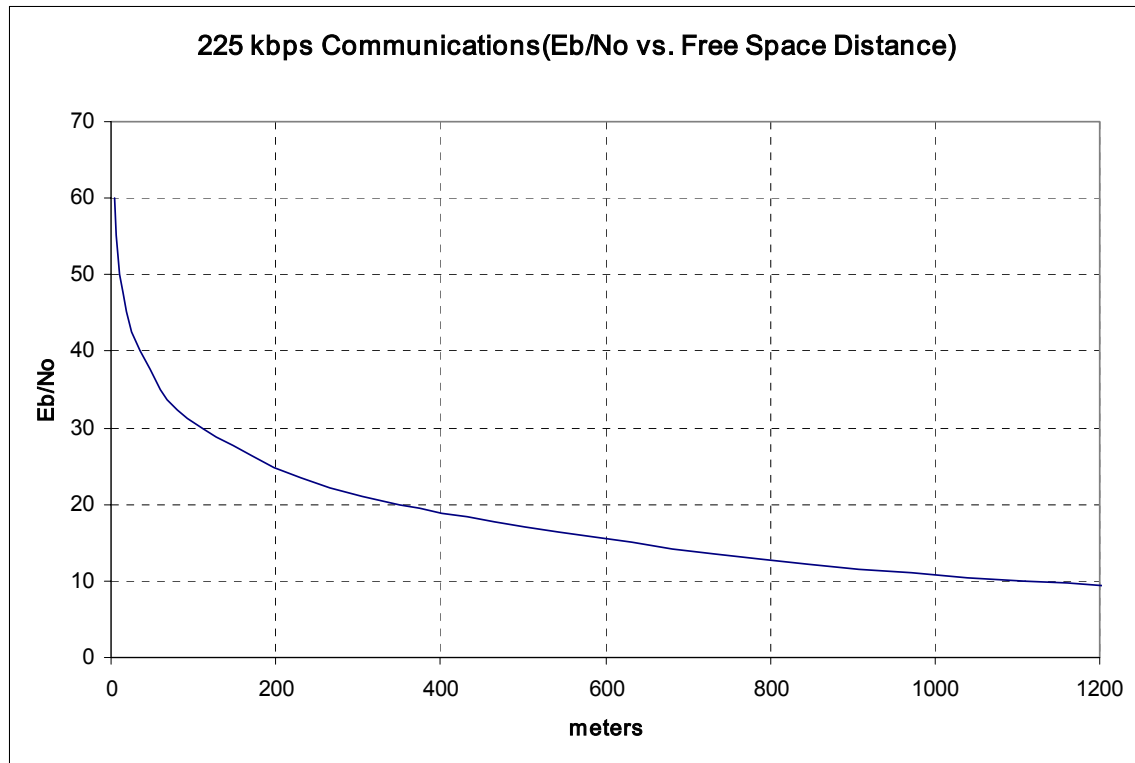


Figure 2 – E_b/N_o vs. Distance for 225 kbps data rate

Assumptions:

- TX power: 25 uW
- Data Rate: 225 kbps
- TX ant: 0 dBi
- RX ant: 0 dBi
- Frequency: 403.5 MHz
- Required E_b/N_o : 10 dB
- Implementation Loss: 3 dB
- RX Noise Figure: 6 dB
- Excess Propagation Loss: 0 dB

The assumption of free space propagation is obviously unrealistic. If we assume an excess propagation loss of 20 dB then we get reduced range for a given signal-to-noise ratio (E_b/N_o) as shown in the plot below.

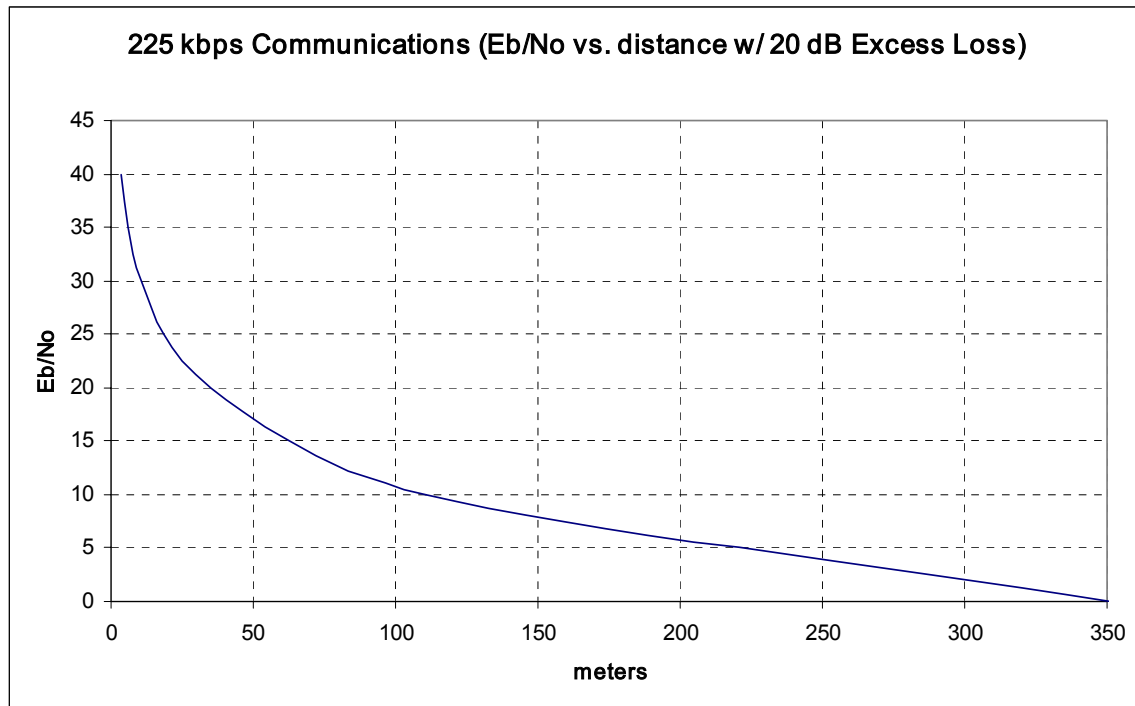


Figure 3 – E_b/N_o vs. Distance for 225 kbps data rate with 20 dB excess propagation loss

Assumptions:

- TX power: 25 uW
- Data Rate: 225 kbps
- TX ant: 0 dBi
- RX ant: 0 dBi

- Frequency: 403.5 MHz
- Required Eb/No: 10 dB
- Implementation Loss: 3 dB
- RX Noise Figure: 6 dB
- Excess Propagation Loss: 20 dB

We can see that even with 20 dB excess propagation loss we are able to maintain an Eb/No > 10 dB out to a range of 100 meters. Thus, it appears that at 400 MHz the 25 uW radio signal goes a considerable distance. Naturally, the reliable communications range will be decreased if the RX antenna efficiency is less than 0 dBi.

Antenna Efficiency Question

We've made the assumption that the TX antenna efficiency is 0 dBi. This assumption is justified due to the fact that the Commission sets the 25 uW limit as EIRP; hence, any TX antenna inefficiency is already taken into consideration.

On the other hand, we've also assumed that the RX antenna efficiency is also 0 dBi. This assumption is probably not valid for the body worn device, but this assumption certainly could be valid for the control device. The justification is that a $\frac{1}{4}$ wavelength at 403.5 MHz is 18.6 cm (7.25 inches) so

that a full size, efficient, vertical antenna using the top of the controller unit as a ground plane is certainly possible; thus, the assumed RX antenna efficiency of 0 dBi seems justified, at least in the case of the controller device.

Assumption on 20 dB excess propagation loss

An issue that is always unknown is “what is the actual excess attenuation of the signal due to propagation within a building”? Or put in other terms, how many dB of signal attenuation does each wall and/or floor in a building cause? There are no fixed answers since it depends entirely on how the wall is constructed (dry wall vs. poured rebar concrete). If we assume 3 dB of attenuation per wall at 400 MHz and 10 dB attenuation per floor then 20 dB excess loss assumes about ½ dozen walls or two floors. Obviously, this is just an educated guess. A web search did not turn up any references that offered a better estimate.

Interference Range – TX signal above thermal noise level

One question is how much interference does the TX signal from the controller cause? In other words - given some assumptions - how far away from the controller transmitter do we need to be before the TX signal falls below the thermal noise floor? The following plots address this issue and give the results in a 1 Hz bandwidth (PSD/Hz) where it was assumed that the TX

power is evenly distributed over 225 kHz of bandwidth. The first plot is for 0 dB excess propagation loss and the second plot is for 20 dB excess propagation loss.

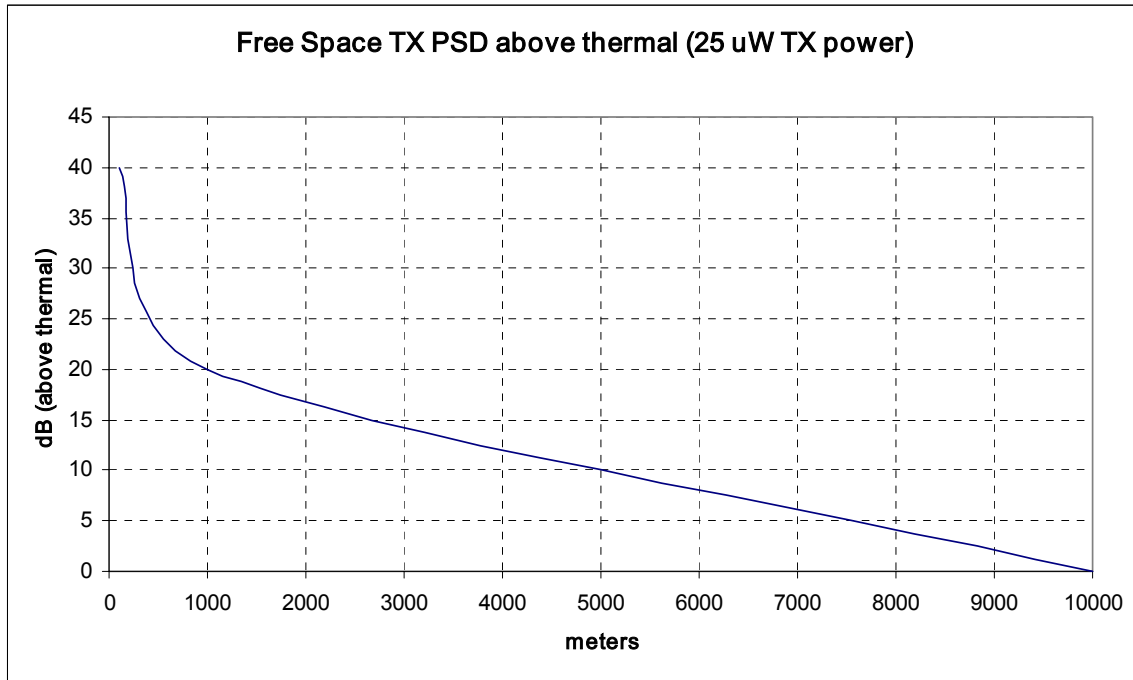


Figure 4 – Free space TX PSD/Hz, dB above thermal

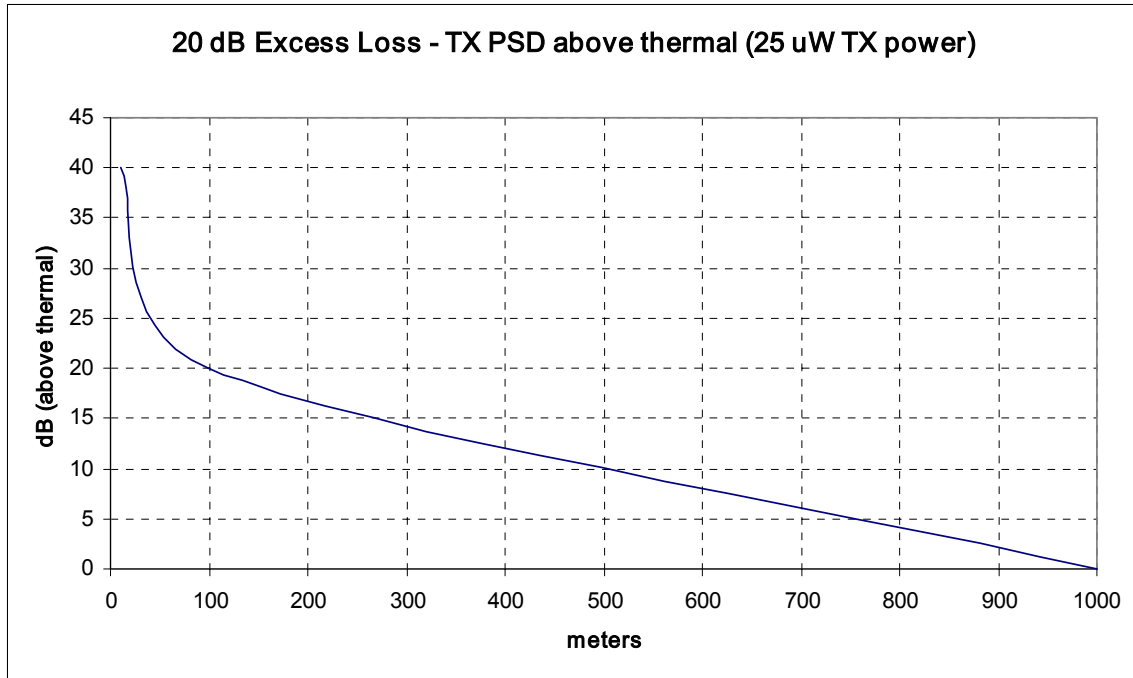


Figure 5 – Free space TX PSD/Hz, dB above thermal, 20 dB excess loss

We can see that even if we assume 20 dB of excess propagation loss the TX signal does not fall below the thermal noise floor until 1 km. This may or may not be a problem depending upon the desired signal of interest field strength at the intended receiving antenna as discussed in Appendix 2.

Appendix 2 (Multi-user Density Analysis – LBT Vs Non-LBT)

Multi-user Density – User separation via physical distance only

As we showed previously, the TX signal doesn't fall below the thermal noise floor until 1 km of distance. However, the amount of interference we experience is dependent upon the SNR between the controller and body unit of interest. Obviously, if the distance between the two units is small then the SNR will be large. For example, from Figure 3 we see that the SNR is 30 dB at 10 meters and 10 dB at 100 meters. So if the normal deployment is at 10 meters or less then we have excess SNR in the link which we can exchange for higher density deployment of units (i.e. can tolerate more interference from other units).

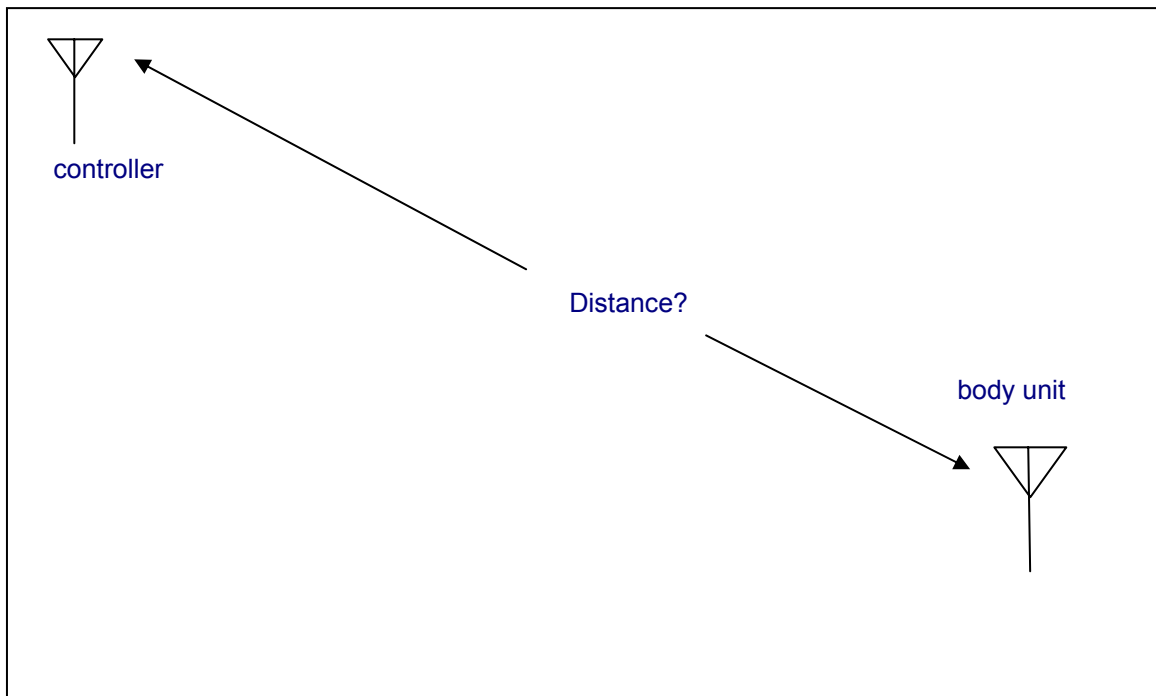


Figure 6 – Deployment scenario - what is distance between controller and
body unit?

Illustrative examples of physical separation at 225 kbps data rate

Case 1 – normal deployment distance is 10 meters

- From figure 3: E_b/N_0 at 10 meters is 30 dB
- The desired E_b/N_0 is 10 dB so we can tolerate ~20 dB above thermal interference
- From figure 5: separation distance for 20 dB of interference above thermal is ~100 meters

Case 2 – normal deployment distance is 100 meters

- From figure 3: E_b/N_0 at 100 meters is 10 dB
- The desired E_b/N_0 is 10 dB so we can tolerate 0 dB above thermal interference
- From figure 5: separation distance for 0 dB of interference above thermal is ~1 km

The bottom line is that relying on physical separation alone to support multiple users is not practical and even a simple protocol would greatly increase the user density. This is discussed in the next several sections.

Multi-user Density – Listen Before Talk (LBT) Protocol

As we showed in the previous section, if we depend solely upon physical distance separation to prevent interference then the deployment density can

become rather sparse. In the original report and order for MICS the Commission instituted a listen-before-talk (LBT) protocol in trying to achieve interference free coexistence between multiple users. In cases where there is excess SNR in the link (short range communications) a LBT protocol will greatly improve coexistence. But as the SNR becomes marginal (e.g. longer distance) the LBT becomes increasingly unreliable.

To explore the relationship between the LBT performance and the number of users, a simulation was written such that each of N users has a 1 second packet of data to send at a random time within a 100 second interval (1% duty cycle). The results below indicate the number of times that the sender had to defer his transmission because the channel was already in use.

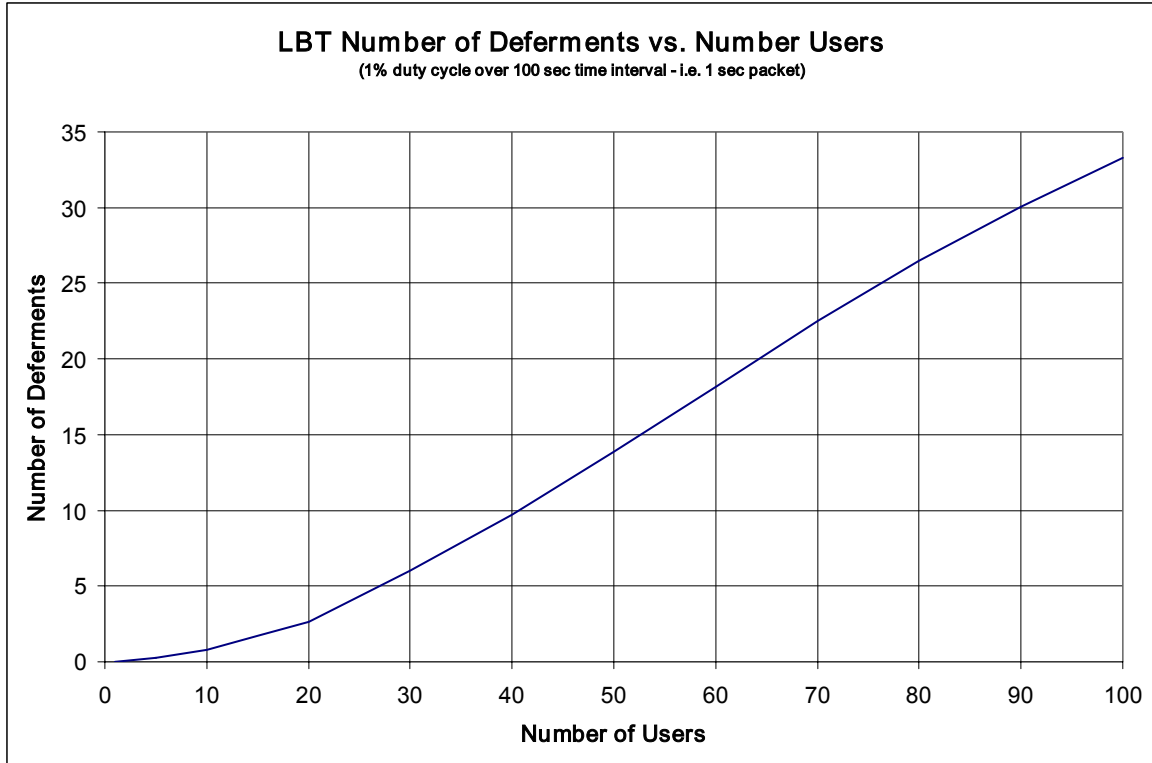


Figure 7 – Number of LBT Deferments vs. Number of Users

In the best case, a transmission deferment does not increase the transmission traffic on the air since a unit will not make a transmission until the LBT protocol indicates that the spectrum is quiet. However, the deferment process does increase the time delay before the message is transmitted. In practice, the LBT protocol *still* has a slight chance of having a packet collision in the case where the media is quiet and two units decide to transmit at the same time. Thus, it would be prudent to include an ACK/NACK handshake (possibility of retransmission) in the protocol to make sure that the information was correctly received.

Non-LBT Protocol with retransmission – an alternative protocol suitable for low duty

An alternative approach to LBT, that is especially effective for use with low duty cycle, is the non-LBT protocol with retransmission. The implementation of this protocol requires that the controller and the body unit be able to do an “acknowledgment” handshake after every data packet is sent. The basic scheme is when a unit has a data packet to send, it just “sends it” without doing LBT. This may or may not result in a collision with another unit’s data packet. After sending the packet, the first unit listens for an acknowledgment. If one is received then communications has successfully occurred. If an acknowledgment is not received (meaning a collision, out of range, offline, ...) then the unit waits a random amount of time (back-off) and retransmits the message (which has the undesirable side affect of increasing the packet traffic on the air). This process continues until the transmission is successfully made. It is highly probable that eventually the message will be successfully sent for the case where all units have a low duty cycle. The question is how much time delay can be tolerated in sending the message.

In order to illustrate the collision rate versus the number of users, a simulation was written such that each of N users sends a 1 second packet of data at a random time within a 100 second interval (1% duty cycle). The results giving the number of collisions are shown below.

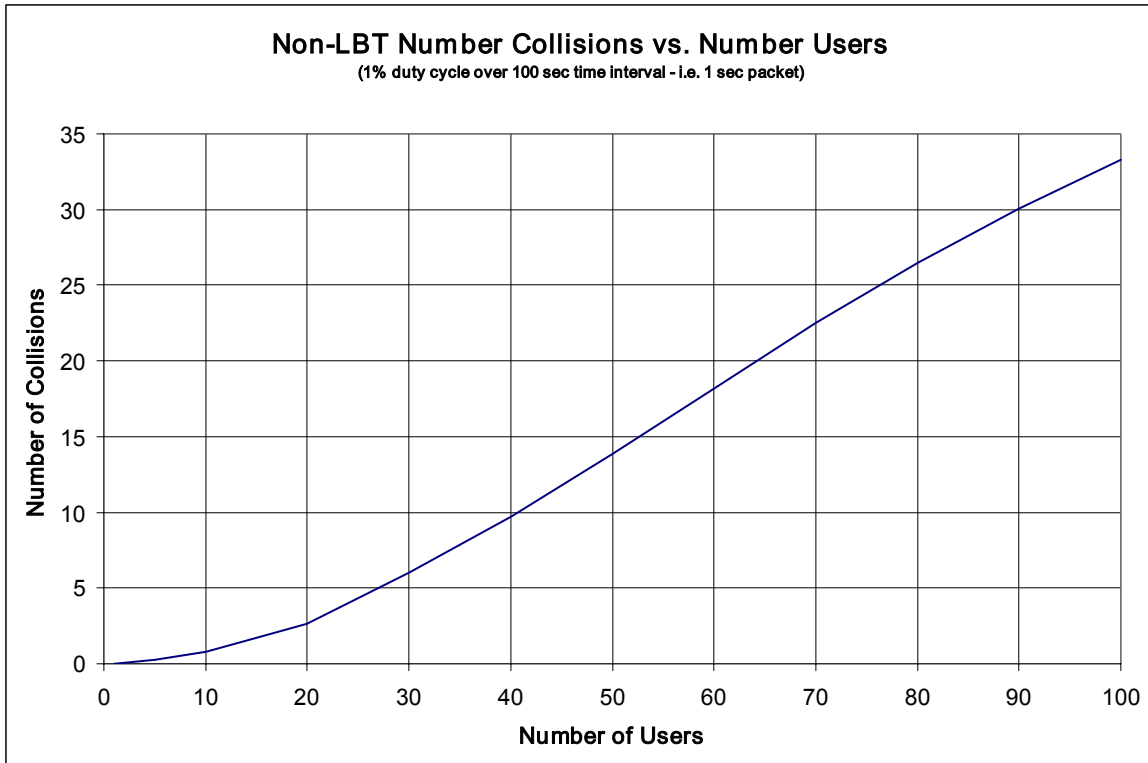


Figure 8 – Number of Non-LBT Collisions vs. Number of Users

One may notice that this is actually the same curve as shown in Figure 7. The difference is in the action taken when a potential collision condition exists. In the case of the non-LBT protocol, we can see that for a small number of users, say 10 users, the number of collisions are relatively small (for $N=10$ the average number of collisions is 0.78). This means that using the non-LBT protocol with retransmission, with $N=10$, on the average ~92% of the time the messages go through without interference. For the 8% of the messages that experience a collision, a retransmission will be required to get the message through. The need for a retransmission will cause a slight delay in

processing the message, which may be acceptable, and will also increase the amount of packet traffic being sent over the air (which increases the chance of a collision in the first place).

The previous figure, Figure 8, showed the number of collisions versus the number of users. If we assume that each collision errors both packets (mutual interference) then the relationship between the number of retransmissions and the number of users, as shown in Figure 9, is two retransmissions for each collision.

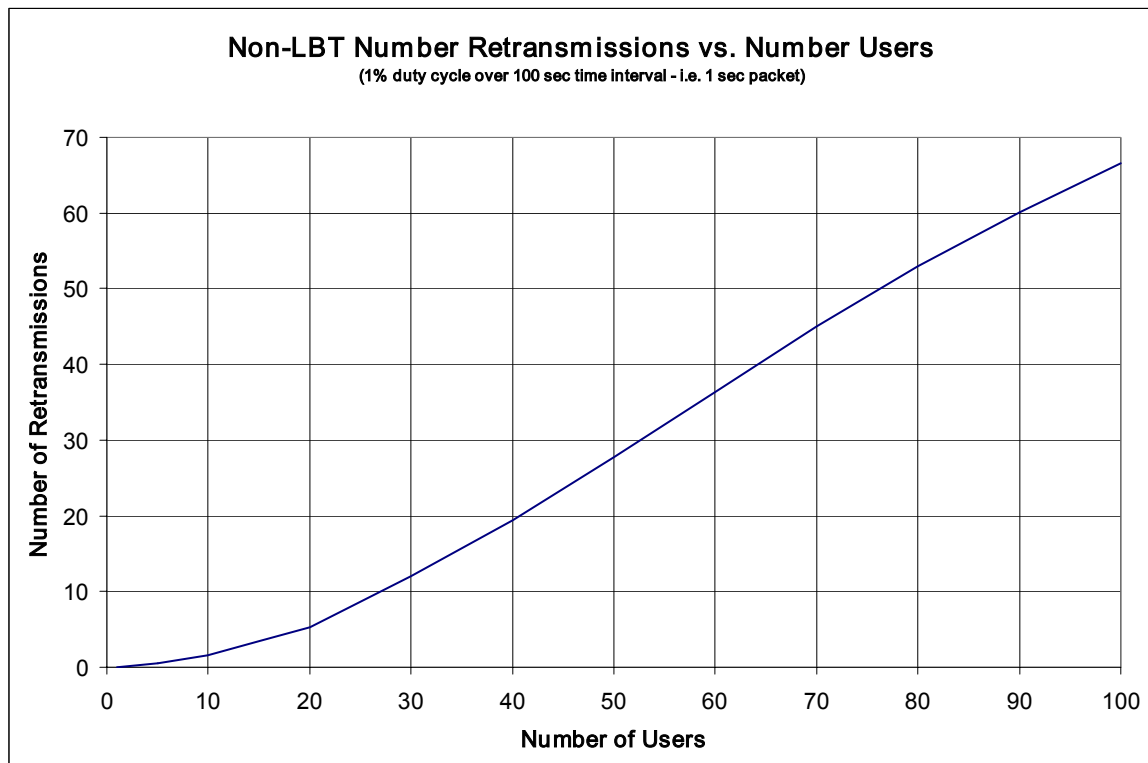


Figure 9 – Number of Non-LBT Retransmissions vs. Number of Users

Keep in mind, the thought behind the non-LBT protocol is that whenever there is a collision between packets, both packets are corrupted. In reality, it would not be unusual for only one packet to be corrupted depending upon the relative signal-to-noise ratios, which would tend to slightly increase the throughput efficiency of the protocol.

Comments on LBT vs. Non-LBT

A natural question would be “which is best, LBT or non-LBT”? As we have discussed, the LBT algorithm is more efficient inasmuch as it minimizes the number of packets transmitted over the air. But we’ve also pointed out that the LBT should include a ACK/NACK retransmission capability to cover those instances when a collisions still occurs. On the other hand, the non-LBT protocol is not as efficient since it relies solely on retransmissions to handle collisions.

Perhaps a terse summary would be that the LBT protocol is more efficient, but it is also more complex (since it has to do a listen before talk function), while the non-LBT protocol is less efficient but also less complex. But it is not clear how this complexity translates to actual hardware cost.

Density based upon either the LBT or Non-LBT Protocol

We use an illustrative example to show an estimate on the user density given some assumptions.

- Data Rate: 225 kbps
- Number of Users: 10
- Deployment Distance: 10 meters

Deployment Area: 314 sq. meters

Density: $314/10 = 31.4$ sq. meters per user

This density would be comparable to a large hospital ward with private rooms. For this given density, it is thought that the throughput of the LBT protocol would be slightly larger than the non-LBT protocol, but at the expense of increased complexity.

Ways to improve user density

- Reduced operating range: the closer the controller and the body unit the better.
- Transmit Power Control: if the distance is short, then the TX power can be reduced below 25 uW so as to reduce unintentional interference while maintaining some minimum SNR.
- Shorter duty cycle: the less time “on-the-air” for each unit, the more units we can support

Comments on Duty Cycle Reduction

The plot below shows the LBT deferments for 0.1% duty cycle. By comparing figure 7 with figure 10, we can see that the number of deferments decreased from 33 deferments for 100 users with 1% duty cycle to only 8 deferments for the same 100 users at 0.1% duty cycle.

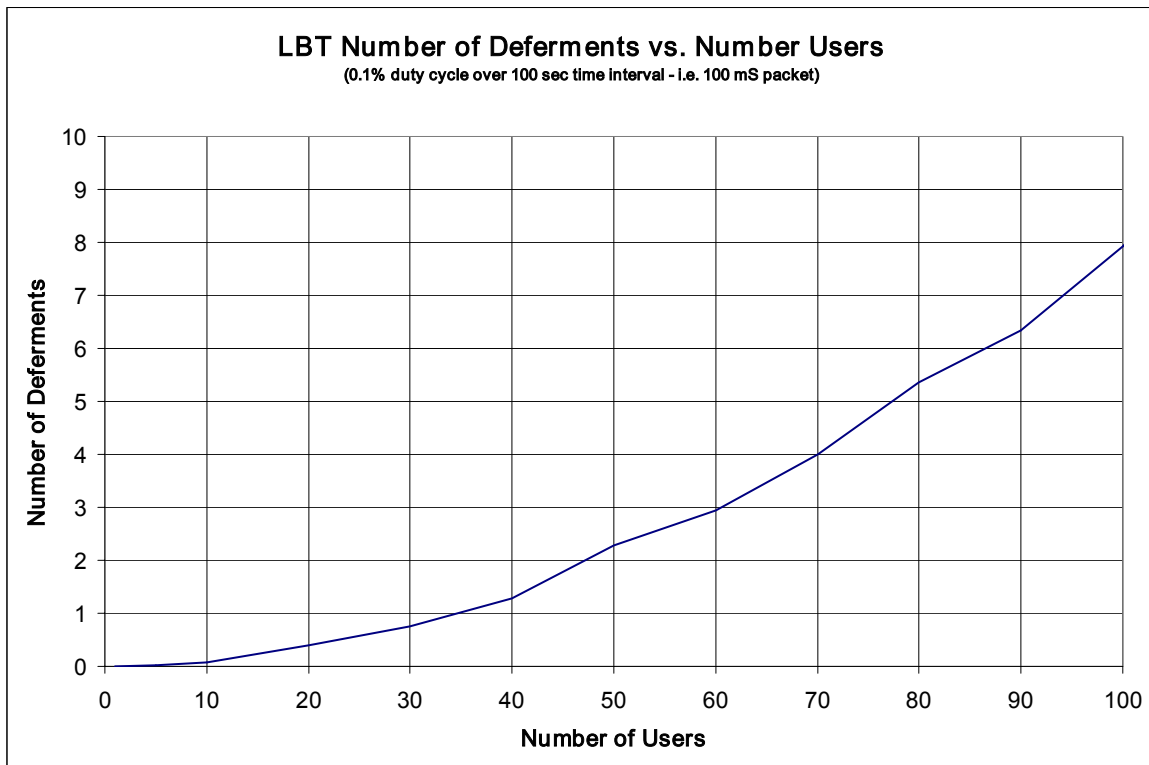


Figure 10 – Number of LBT Deferments vs. Number of Users for 0.1% duty cycle

Appendix 3 (Trading-off Increased BW for Reduced Duty Cycle)

The ability to arbitrarily reduce the duty cycle is limited by the amount of data that needs to be sent given the packet time duration. While it is unknown how much data needs to be sent per unit time, one potential design trade-off is changing the modulation bandwidth to control the duty cycle; that is, reduce the duty cycle by sending the data burst faster. The next plot illustrates this point with the following assumptions.

Analysis Assumptions for Illustrative Example

Operating Distance: 10 meters

Bits per modulation symbol: 1 bit/symbol (e.g. BPSK)

Bits per packet: 225 kbits

The following figure shows how we can trade-off the symbol rate against duty cycle.

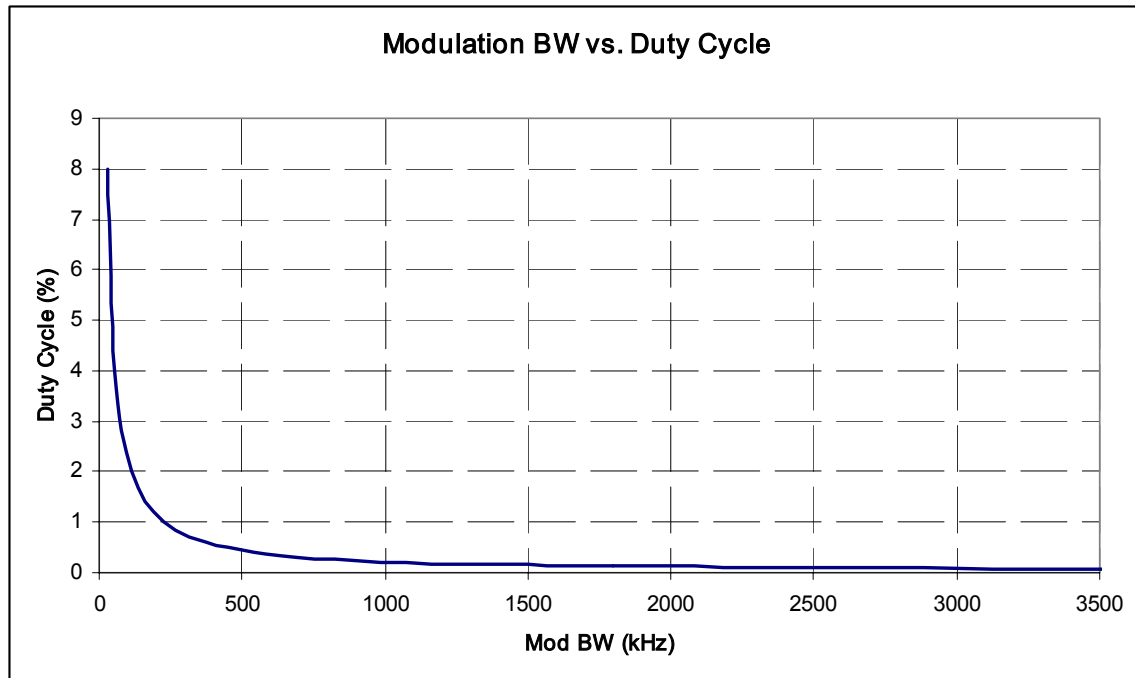


Figure 11 – A wider Modulation BW gives a lower Duty Cycle for a given packet length

Naturally, the wider bandwidth (higher bit rate) data packet burst will require more signal to noise ratio (E_b/N_0) for transmission. The next figure shows the modulation bandwidth versus the actual E_b/N_0 given the following assumptions:

- Operating Distance: 10 meters
- TX power: 25 μ W
- Modulation BW: variable
- Modulation Efficiency: 1 bps per Hertz
- TX ant: 0 dBi

- RX ant: 0 dBi
- Frequency: 403.5 MHz
- Required Eb/No: 10 dB
- Implementation Loss: 3 dB
- RX Noise Figure: 6 dB
- Excess Propagation Loss: 0 dB

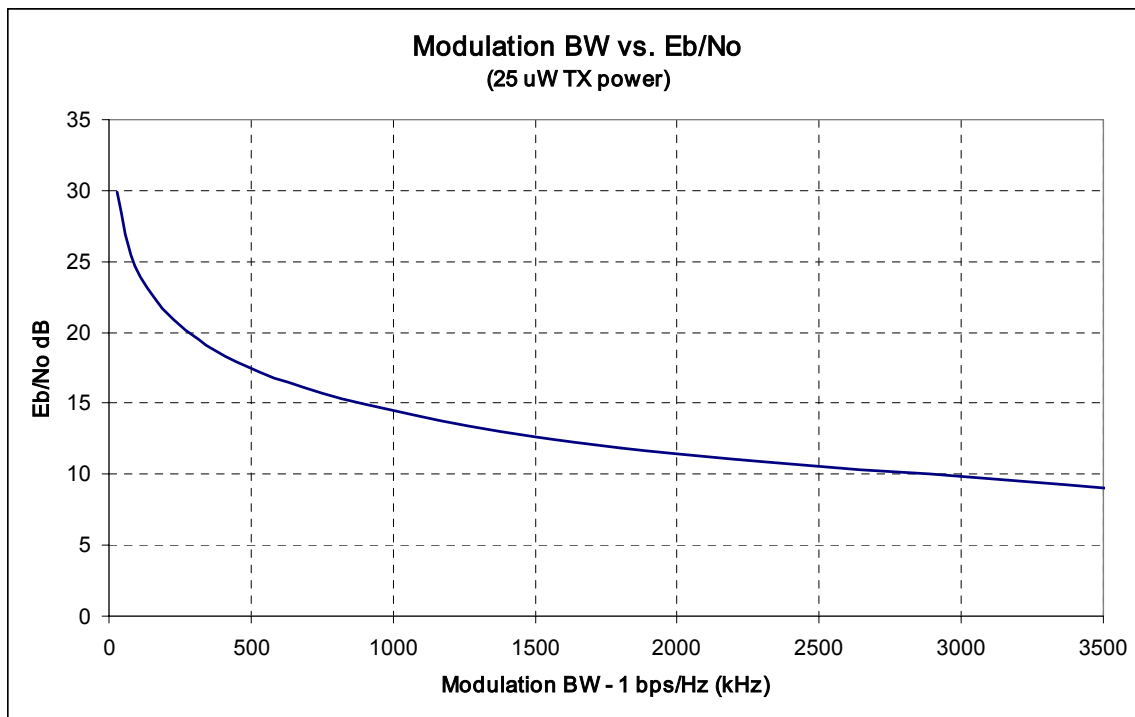


Figure 12 – For the above assumptions - Modulation BW vs. Eb/No at 10 meters

We stated that we required at least 10 dB of Eb/No for satisfactory communications; thus, the allowed modulation bandwidth at 10 meters range,

for the given assumptions, is < 3 MHz. It turns out that a modulation bandwidth slightly less than 3 MHz, for the given assumptions of figure 11, yields a duty cycle of 0.1%. Referring to figure 10, a duty cycle of 0.1% and $N=100$ yields an 8% collision rate. We can compare this to the results shown in the paragraph following Figure 8 where it was shown that with a 1% duty cycle, $N=10$ users gave us an 8% collision rate. To summarize: by increasing the modulation bandwidth by a factor of 10, we reduced the duty cycle by a factor of 10, which in turned allowed us to accommodate 10x the number of users while maintaining the same packet collision rate.

Appendix 4 (Decreased transmit power analysis)

In support of analyzing 250 nW vs. 25 uW for use in the wing bands without LBT, we developed two informative charts. Figure 13 shows the relationship between communication range, given the stated assumptions, and the TX power. Figure 14 shows the relationship between the user density, for the stated assumptions, and the TX power. We are particularly interested in the TX power levels at 250 nW and 25 uW. The RX antenna efficiency of -20 dBi is an estimate for a body worn device.

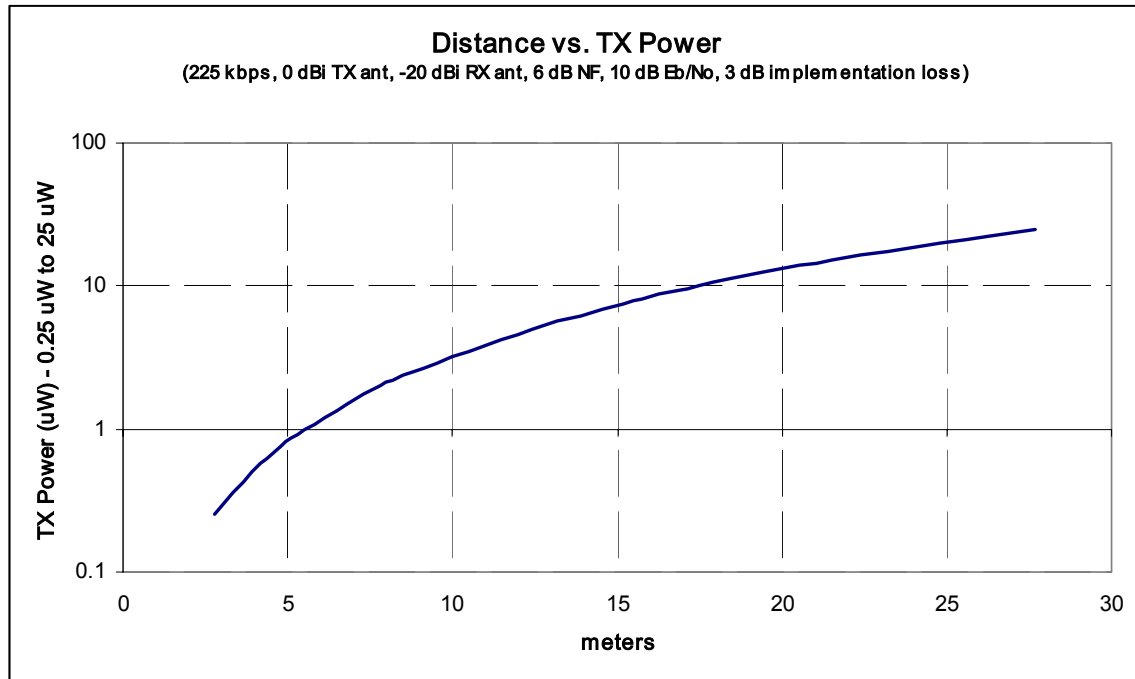


Figure 13 – Distance vs. TX Power ... of interest is 250 nW and 25 uW

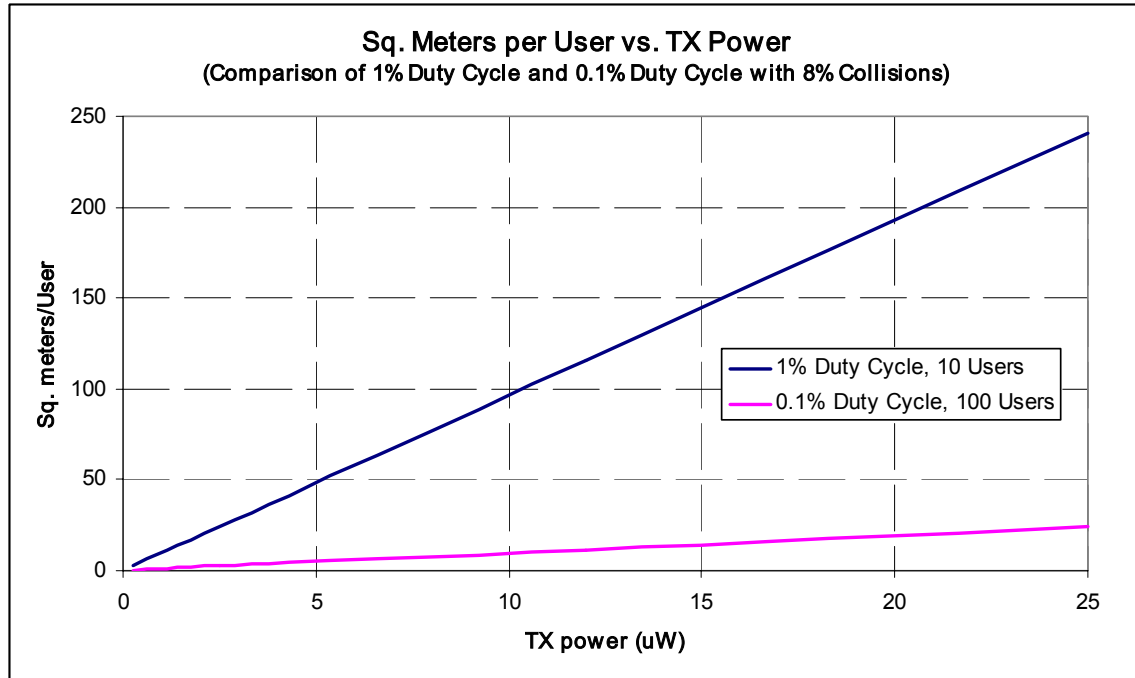


Figure 14 – User density vs. TX power for two duty cycles (1% and 0.1%)

In Figure 14, the measure of “goodness” is having a small sq. meter/ user number; that is, the smaller the square meter/user number, the larger the number of users we can insert into a given area. We can see from Figure 14 that two factors influence the user density: the TX power and the Duty Cycle. We can see that the graph in Figure 14 becomes relatively insensitive to the TX power as the duty cycle is decreased; thus, it is possible to use a higher TX power higher and still maintain good user density if the duty cycle is reduced. We’d suggest the NPRM text allow full power at 0.1% duty cycle and then proportionally reduce the TX power as the duty cycle increases. This concept is illustrated in the table below.

Duty Cycle	TX Power
0.1 %	25.0 uW
1.0 %	2.5 uW

10.0 %	0.25
--------	------